●                    ●

WHAT IS CLAIMED IS:

1   1.  In a switch having $N_{in}$ input ports applied to $K_{in}$ input

2       shared blocks, a central switching fabric, and $N_{out}$

3       output ports provided from $K_{out}$ output shared blocks, a

4       method for scheduling packets queued at the input

5       shared blocks for application to the output ports, the

6       method comprising steps of:

7       a)   providing, for each of the input shared blocks, an

8       indication of a number of links reserved by the input

9       shared block to each of the output shared blocks;

10      b)   providing each of the input shared blocks with a

11     token, each token corresponding to an output shared

12     block and including a value indicating a number of

13     links available to the associated output shared block;

14     c)   if it is determined that an input shared block

15    needs links to an output shared block associated with

16    a token held by the input shared block, then

17          i)   reserving links, to the extent available as

18          indicated by the token, to the output shared

19          block

20          ii)   updating the value indicating the number of

21          links available to the associated output shared

22          block, and

23          iii)   updating the value of the indication of a

24          number of links reserved by the input shared

25          block to the associated output shared block.

1   2.  The method of claim 1 further comprising a step of

2       d)   passing the token to a next input shared block at

3       the end of a reservation time slot.

1   3.   The method of claim 2 wherein a cell slot includes at

2       least one reservation slot and further comprising a

3       step of:

4         d)    delivering cells to the central switch fabric

5       based on their currently reserved links at the end of

6       each cell slot.


1   4.   The method of claim 3 further comprising a step of:

2         e)   destroying the tokens at the end of each cell

3       slot; and

4         f)   generating new tokens at each input shared block

5       at the end of each cell slot.


1   5.   In a switch having $N_{in}$ input ports applied to $K_{in}$ input

2   shared blocks, a central switching fabric, and $N_{out}$ output

3   ports provided from $K_{out}$ output shared blocks, a method for

4   scheduling packets queued at the input shared blocks for

5   application to the output ports, the method comprising

6   steps of:

7         a)    for each of the input shared blocks, providing a

8       request token associated with one of the output shared

9       blocks, each of the request tokens including an

10      indication based on a number of requested links for

11      the output shared block with which it is associated;

12      b)   for each of the input shared blocks, providing a

13      release token associated with one of the output shared

14      blocks, each of the release tokens including an

15      indication based on a number of released links for the

16      output shared block with which it is associated;

17      c)   for each of one or more reservation time slots

18      within a cell time slot,

19  i)  accepting, by an input shared block, a

20  request token from another input shared block,

21  ii)  determining whether a virtual output queue

22  of the input shared block associated with the

23  output shared block with which the request token

24  is associated, is heavily occupied,

25  iii)  if it is determined that the virtual output

26  queue of the input shared block associated with

27  the output shared block with which the request

28  token is associated is heavily occupied, then

29      A)  requesting at least one extra link to

30      the output shared block associated with the

31      accepted request token,

32  iv)  determining whether the virtual output queue

33  of the input shared block associated with the

34  output shared block with which the request token

35  is associated, is lightly occupied,

36  v)  if it is determined that the virtual output

37  queue of the input shared block associated with

38  the output shared block with which the request

39  token is associated is lightly occupied, then

40      A)  releasing at least one link to the

41      output shared block associated with the

42      accepted request token if it is indicated

43      that a number of requested links for the

44      output shared block is greater than zero,

45  vi)  if it is determined that the virtual output

46  queue of the input shared block associated with

47  the output shared block with which the request

48  token is associated is not lightly occupied, then

49      A)  releasing at least one link to the

50      output shared block associated with the

51                     accepted request token if the input shared

52                     block reserved more than a predetermined

53                     number of links and if it is indicated that

54                     a number of requested links for the output

55                     shared block is greater than zero,

56           vii)   accepting, by the input shared block, a

57           release token from another input shared block;

58           viii)   determining whether or not the input

59           shared block can take a link to the output shared

60           block associated with the release token,

61           ix)   if it is determined that the input shared

62           block can take a link to the output shared block

63           associated with the release token, then taking a

64           link from the release token.


1    6.    The method of claim 5 wherein the step of determining

2    whether a virtual output queue of the input shared block

3    associated with the output shared block with which the

4    request token is associated, is heavily occupied, is based

5    on a comparison with a threshold value.


1    7.    The method of claim 5 wherein the step of requesting an

2    extra link to the output shared block associated with the

3    accepted request token, is effected by setting a request

4    indicator corresponding to the input shared block and the

5    output shared block, and incrementing the indication based

6    on the number of links to the output shared block

7    requested.


1    8.    The method of claim 5 wherein the step of determining

2    whether the virtual output queue of the input shared block

3    associated with the output shared block with which the

4 request token is associated, is lightly occupied, is based

5 on a comparison with a threshold value.

1 9. The method of claim 5 wherein the step of releasing a

2 link to the output shared block associated with the

3 accepted request token is effected by decreasing the

4 indication based on the number of links to the output

5 shared block released.

1 10. The method of claim 5 wherein the step of determining

2 whether or not the input shared block can take a link to

3 the output shared block associated with the release token,

4 is based on a number of all reserved links by the input

5 shared block and an indication of whether or not the input

6 shared block had requested a link to the output shared

7 block.

1 11. The method of claim 5 further comprising a step of:

2     d) delivering, by each of the input shared blocks,

3        cells to the central switch fabric based on current

4        indications of a number of link reservations to each

5        of the output shared blocks, at the end of a cell time

6        slot.

1 12. In a switch having $N_{in}$ input ports applied to $K_{in}$ input

2 shared blocks, a central switching fabric, and $N_{out}$ output

3 ports provided from $K_{out}$ output shared blocks, a method for

4 scheduling packets queued at the input shared blocks for

5 application to the output ports, the method comprising

6 steps of:

7     a)    for each of the input shared blocks, providing a

8           request token associated with one of the output shared

9     blocks, each of the request tokens including an

10    indication based on a number of requested links for

11    the output shared block with which it is associated;

12    b)    for each of the input shared blocks, providing a

13    release token associated with one of the output shared

14    blocks, each of the release tokens including an

15    indication based on a number of released links for the

16    output shared block with which it is associated;

17    c)    for each of one or more reservation time slots

18    within a cell time slot,

19       i)    accepting, by an input shared block, a

20       request token from an another input shared block,

21       ii)    determining whether a virtual output queue

22       of the input shared block associated with the

23       output shared block with which the request token

24       is associated, is heavily occupied,

25       iii)    if it is determined that the virtual output

26       queue of the input shared block associated with

27       the output shared block with which the request

28       token is associated is heavily occupied, then

29         A)    requesting at least one extra link to

30         the output shared block associated with the

31         accepted request token,

32       iv)    determining whether the virtual output queue

33       of the input shared block associated with the

34       output shared block with which the request token

35       is associated, is lightly occupied,

36       v)    if it is determined that the virtual output

37       queue of the input shared block associated with

38       the output shared block with which the request

39       token is associated is lightly occupied, then

A)    releasing at least one link to the
output shared block associated with the
accepted request token if it is indicated
that a number of requested links for the
output shared block is greater than zero,

vi)    if it is determined that the virtual output
queue of the input shared block associated with
the output shared block with which the request
token is associated is not lightly occupied, then

A)    releasing at least one link to the
output shared block associated with the
accepted request token if the input shared
block reserved more than a predetermined
number of links,

vii)    accepting, by the input shared block, a
release token from another input shared block,

viii)    determining whether to release at least
one link to the output shared block with which
the accepted release token is associated based on
a queue occupancy, a number of links reserved,
and a predetermined number of links

ix)    if it has been determined to release a link
to the output shared block with which the
accepted release token is associated, releasing a
link,

x)    determining whether or not to take at least
one released link to the output shared block with
which the accepted release token is associated
based on queue occupancy, a number of links
reserved, and a number of links between the input
shared block and the central switch fabric, and

71      xi)   if it is determined to take at least one
72      released link to the output shared block with
73      which the accepted release token is associated,
74      taking at least one link.

1   13.   The method of claim 12 wherein the step of determining
2   whether a virtual output queue of the input shared block
3   associated with the output shared block with which the
4   request token is associated, is heavily occupied, is based
5   on a comparison with a threshold value.

1   14.   The method of claim 12 wherein the step of requesting
2   at least one extra link to the output shared block
3   associated with the accepted request token, is effected by
4   setting a request indicator corresponding to the input
5   shared block and the output shared block, and incrementing
6   the indication based on the number of links to the output
7   shared block requested.

1   15.   The method of claim 12 wherein the step of determining
2   whether the virtual output queue of the input shared block
3   associated with the output shared block with which the
4   request token is associated, is lightly occupied, is based
5   on a comparison with a threshold value.

1   16.   The method of claim 12 wherein the step of releasing
2   at least one link to the output shared block associated
3   with the accepted request token is effected by decreasing
4   the indication based on the number of links to the output
5   shared block released.

1  17.  The method of claim 12 wherein the step of determining

2  whether or not the input shared block can take at least one

3  link to the output shared block associated with the release

4  token, is based on a queue occupancy of a virtual output

5  queue, and a number of all reserved links by the input

6  shared block and an indication of whether or not the input

7  shared block had requested at least one link to the output

8  shared block if it is indicated that a number of released

9  links for the output shared block is greater than zero.


1  18.  The method of claim 12 further comprising a step of:

2       d) delivering, by each of the input shared blocks,

3          cells to the central switch fabric base based on

4          current indications of a number of link reservations

5          to each of the output shared blocks, at the end of a

6          cell time slot.


1  19.  A switch for switching packets arriving at a number of

2  input ports to an appropriate one of a number of output

3  ports, the switch comprising:

4       a)  a central switching fabric;

5       b)  output shared blocks, each coupled with at least

6          one output port;

7       c)  links between the central switch fabric and each

8          of the output shared blocks

9       d)  input shared blocks, each

10             i)  coupled with at least one input port,

11             ii)  having virtual output queues, each of the

12                virtual output queues corresponding to one or

13                more output ports,
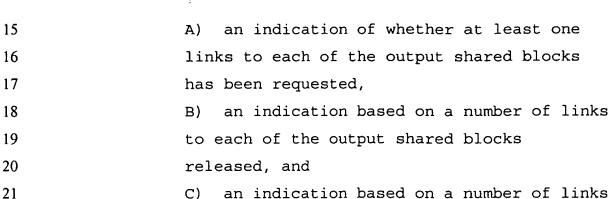
14             iii)  storing

84

15          A)    an indication of whether at least one

16          links to each of the output shared blocks

17          has been requested,

18          B)    an indication based on a number of links

19          to each of the output shared blocks

20          released, and

21          C)    an indication based on a number of links

22          reserved to each of the output shared

23          blocks;

24    e)    request tokens, each associated with a particular

25    one of the output shared blocks and each indicating a

26    number of requests for links to the associated one of

27    the output shared blocks;

28    f)    release tokens, each associated with a particular

29    one of the output shared blocks and each indicating a

30    number of released links to the associated one of the

31    output shared blocks; and

32    g)    links between the central switch fabric and each

33    of the input shared blocks.

1    20.   The switch of claim 19 wherein each of the input

2    shared blocks holds at least one of the request tokens and

3    at least one of the release tokens during a reservation

4    time slot.